

## CLAIMS

What is claimed is:

- 1 1. A method of accessing shared memory in a computer system having a  
2 plurality of nodes, including a first node, wherein each node includes a processor and  
3 local memory, the method comprising:  
4 distributing an application across the plurality of nodes;  
5 building an application virtual address space, wherein building an application  
6 virtual address space includes:  
7 building a local virtual address space for the application in each of the  
8 plurality of nodes, wherein the local virtual address space translates a virtual  
9 address generated by the application executing on that node to a physical  
10 address in local memory for that node; and  
11 exporting the local virtual address space for each node to a Remote  
12 Translation Table (RTT) associated with that node; and  
13 performing a memory reference to a memory location in the application  
14 virtual address space, wherein performing a memory reference to a memory location  
15 in the application virtual address space includes translating bits of the application  
16 virtual address into a node address associated with the first node and translating bits  
17 of the application virtual address using the RTT associated with the first node.  
18
- 1 2. The method of claim 1, wherein the local address space is read from a  
2 Translation Look-aside Buffer (TLB).
- 1 3. The method of claim 1, wherein building an application virtual address space  
2 further includes performing a synchronization operation that causes at least some of  
3 the plurality of nodes to wait for all nodes to complete exporting their respective  
4 local virtual address spaces.

1     4.     A system comprising:  
2             a plurality of nodes, each node including:  
3                 one or more processors;  
4                 a memory; and  
5                 a memory controller operatively coupled to the memory and the one  
6             or more processors, wherein the memory controller includes a Remote  
7             Translation Table (RTT), wherein the RTT translates a virtual address  
8             received as part of a memory request received from another node into a  
9             memory request with physical addresses into the memory on the node  
10            associated with the RTT;  
11            further wherein the RTT is initialized upon the start of a process  
12            associated with an application by building virtual to physical address  
13            translations for local virtual address space in the node corresponding to the  
14            application, and exporting the virtual to physical address translations for the  
15            local virtual address space from the node to the Remote Translation Table  
16            (RTT) associated with that node.

1     5.     The system of claim 4, wherein each of the plurality of nodes executes a  
2             synchronization operation that causes at least some of the plurality of nodes to wait  
3             for all of the plurality of nodes to complete exporting the virtual to physical address  
4             translations to their respective Remote Translation Tables.

1     6.     A device-readable medium having instructions thereon that, when executed  
2             on a properly programmed information-processing device having a plurality of  
3             nodes, including a first node, each node having one or more processors, a memory,  
4             and a memory controller and coupled to the memory and the one or more processors,  
5             causes the information-processing device to perform a method comprising:  
6                 distributing an application across the plurality of nodes;  
7                 building an application virtual address space, wherein building an application  
8             virtual address space includes:

9 building a local virtual address space for the application in each of the  
10 plurality of nodes, wherein the local virtual address space translates a virtual  
11 address generated by the application executing on that node to a physical  
12 address in local memory for that node; and  
13 exporting the local virtual address space for each node to a Remote  
14 Translation Table (RTT) associated with that node; and  
15 performing a memory reference to a memory location in the application  
16 virtual address space, wherein performing a memory reference to a memory location  
17 in the application virtual address space includes translating bits of the application  
18 virtual address into a node address associated with the first node and translating bits  
19 of the application virtual address using the RTT associated with the first node.

1 7. The device-readable medium of claim 6, wherein building a local virtual  
2 address space further includes performing a synchronization operation that causes at  
3 least some of the plurality of nodes to wait for all nodes complete exporting their  
4 respective address space.

1 8. The device-readable medium of claim 6, wherein the local address space is  
2 read from a Translation Look-aside Buffer (TLB).

1 9. A multinode system for implementing remote address translation, the system  
2 comprising:

3 a plurality of nodes, including a first node, each of the plurality of nodes  
4 including:  
5 one or more processors,  
6 a memory, and  
7 a memory controller operatively coupled to the memory and the one  
8 or more processors;  
9 means for distributing an application across the plurality of nodes;  
10 means for building an application virtual address space, wherein the means

11 for building an application virtual address space includes:  
12 means for building a local virtual address space for the application in  
13 each of the plurality of nodes, wherein the local virtual address space  
14 translates a virtual address generated by the application executing on that  
15 node to a physical address in local memory for that node; and  
16 means for exporting the local virtual address space for each node to a  
17 Remote Translation Table (RTT) associated with that node; and  
18 means for performing a memory reference to a memory location in the  
19 application virtual address space, wherein performing a memory reference to a  
20 memory location in the application virtual address space includes:  
21 means for translating bits of the application virtual address into a node  
22 address associated with the first node, and  
23 means for translating bits of the application virtual address using the  
24 RTT associated with the first node.

1 10. The multinode system of claim 9, wherein building an application virtual  
2 address space further includes a means for performing a synchronization operation  
3 that causes at least some of the plurality of nodes to wait for all nodes to complete  
4 exporting their respective local virtual address spaces.

1 11. A multi-node system for implementing remote address translation, the system  
2 comprising:  
3 a network;  
4 a source node coupled to the network, wherein the source node includes a  
5 first remote-translation table (RTT);  
6 a remote node coupled to the network, wherein the remote node includes a  
7 second RTT;  
8 wherein on the remote node the second RTT is built using a first local  
9 address space on the source node exported from the source node to the remote node;  
10 and

- 11            wherein on the source node the first RTT is built using a second local address
- 12   space on the remote node exported from the remote node to the source node.